

Knowledge Graphs

AI THAT WORKS LIKE HUMANS THINK

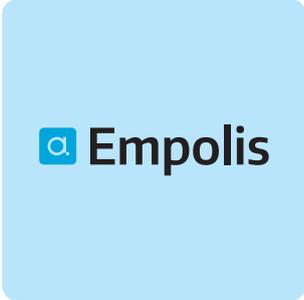


Table of Contents

Management Summary	1
What are knowledge graphs and how do they work?	2
Areas of application	3
• Intelligent content management	3.1
• Intelligent product information	3.2
• Knowledge graphs for service	3.3
Capabilities of knowledge graphs	4
How are knowledge graphs created?	5
Summary	6



Management Summary

Knowledge graphs are a new form of artificial intelligence that work the way humans think. This makes knowledge graphs well suited for widespread use in business, and managers in large and medium-sized companies and organizations are increasingly discovering their value as part of the digital transformation for business and administration.

The high relevance of this technology was repeatedly confirmed by the Gartner Hype Cycle for artificial intelligence. The capacity of a company to use data specifically as its capital for its decision-making and its own business will separate the winners from the losers. Intelligent data analysis and interpretation allow companies to better understand their customers and form the basis for new, profitable business models. At the same time,

data serves as the foundation for innovations and more efficient management of a company's own business processes and resources. Many companies, however, have large quantities of data that are mostly heterogeneous, distributed across multiple locations and departments, and no longer manageable. Organizational data silos and interface issues also make it difficult for companies to utilize their data effectively.

Although there are various AI-based approaches to addressing the issue, they often lack transparency and are too complex for people to understand. In many cases, entire teams of experts are needed to reconstruct and understand the complex algorithms involved – which disqualifies them from widespread commercial use.



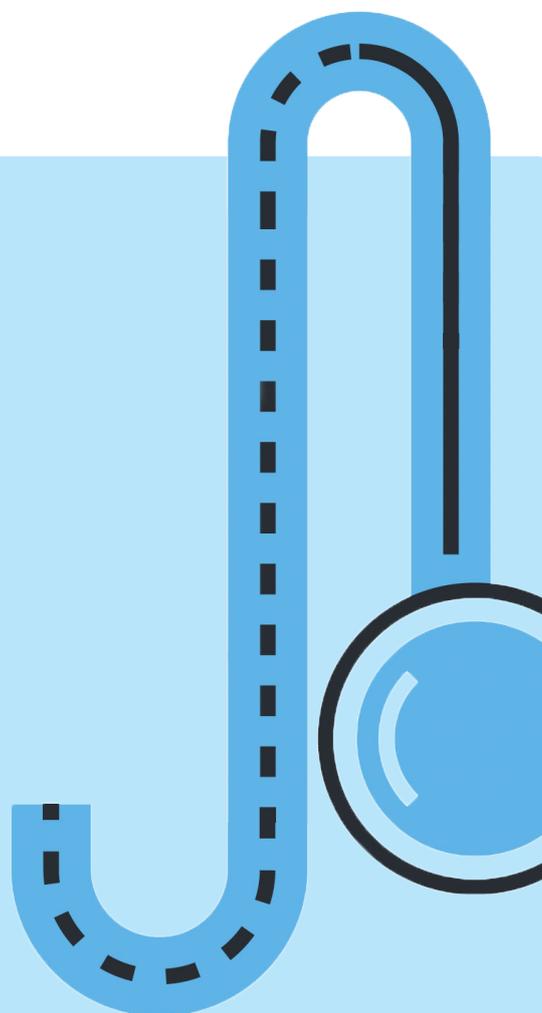


Knowledge graph technology, on the other hand, is a form of artificial intelligence that can be understood and explained. It is a technology that allows users to provide sound answers to complex domain-specific questions by representing subject areas and business processes in terms of their relationships.

Knowledge graphs create links and relationships, whether they be simple and uniform or complex and varied, and enable models that can be used by both humans and computers. By providing the basis for such models, knowledge graphs take the crucial step from data storage to knowledge representation of the kind required for many modern AI processes, and

this ultimately allows users to efficiently monetize their own data capital.

This whitepaper explains in detail how knowledge graphs work and how they can contribute to successful business outcomes. It provides an in-depth view of how knowledge graph technology makes distributed knowledge available at the right time and in the right context. On this basis, fields like service, sales, after-sales, and technical documentation can be revolutionized and transformed from cost centers into profit centers.



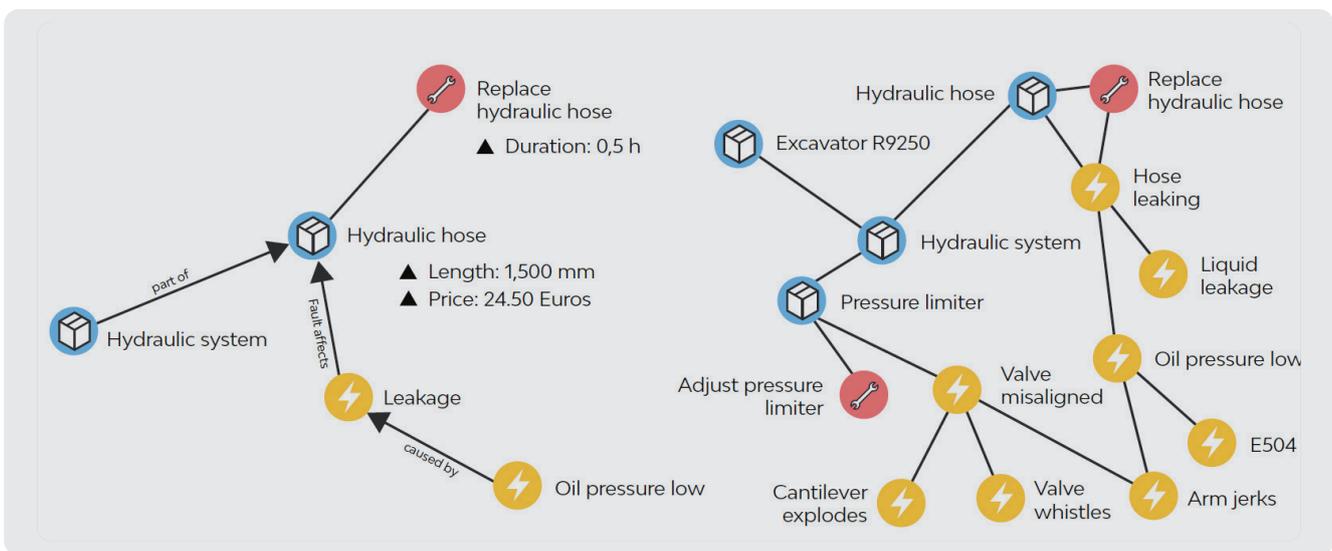


WHAT ARE KNOWLEDGE GRAPHS AND HOW DO THEY WORK?

Knowledge graphs represent knowledge as connections between objects, whereas both objects and connections have different types (also called “relations”) and play the essential role in the matter. In the example in Figure 1, objects of type “Fault” are linked with objects of type “Component” by the relation “Occurs With.” Components are “Part Of” other compo-

nents or of products, and they can be involved in not only faults but also maintenance activities, for instance.

Furthermore, objects can have attributes. In the example shown, components have dimensions and a price; maintenance activities have a duration.



There are also mechanisms to derive knowledge, infer conclusions, and inherit information. For example, the hydraulics system is a part of Excavator R9250, and the pressure limiter is part of the hydraulics system. Therefore, the pressure limiter is also part of the Excavator R9250.

This means that knowledge graphs do more than simply receive data and display it back. They represent the logic of some domain – machines, compliance regulations, product portfolios, or industrial facilities – in a formal model, which can do things like:

Answer questions by evaluating the connections in the graph:

What faults can occur in a model R9250 excavator? And what faults occur only in components that are not installed in the model R9250?

If we observe a jerking of the arm, the knowledge graph can identify a misaligned valve or a leaking hydraulic hose as possible causes, and it can tell us how to determine whether either of those problems is present.



Find similar cases

Given an existing problem, we can use stored knowledge in the knowledge graph to find past cases with similar symptoms in similar components (possibly even in similar usage scenarios).

Examine hypothetical situations

The graph can identify the effects of a fault: Do we have to stop the machine and interrupt work in order to adjust the pressure limiter? What is the worst-case scenario? In the example, all we have to do is follow the cause-and-effect relations to arrive at serious damage. The knowledge graph can also take into account all the possible parameters of a given situation, such as the frequency of damage in the past, weights, forces, expected heat generation, etc., in order to identify the probability and severity of possible damage.

The graph can accommodate any n:m relations (where the relation exists between any number of objects on either side) and provides for inferences as well as inheritance. This makes handling variants and exceptions extremely easy. This includes faults that occur only in a certain variant of a component or only in components with certain features (e.g., hydraulic hoses made of PVC) or with certain combinations of components (e.g., the pressure limiting valve often becomes misaligned only in conjunction with control X).

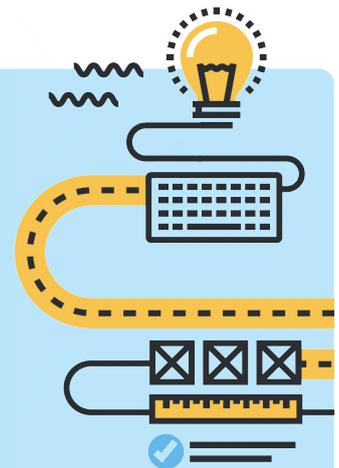
In contrast to other information technologies, knowledge graphs also have a range of properties that are especially useful for digital models:

In a knowledge graph, each logical object exists only once in the data structure. Data are organized around the objects and not according to a rigid arrangement of rows and columns as in a relational database. There is no need to establish object identity or normalization through artificial means, which would introduce added technical complexity involving foreign keys, for example.

Knowledge graphs form comprehensible models, and they can clearly justify their inferences and recommendations at any time (e.g., similar cases and possible causes). This sets them apart from the typical black box models in machine learning. Instead of being trained solely with a large quantity of sample data, they represent explicit expert knowledge.

Graphs can be viewed from all sides. It is just as easy to focus on a component as it is to focus on a fault or a maintenance activity. Starting from a product, we can list all of the possible faults. Conversely, given a fault, we can list all of the components in which it can occur and all of the products that use these components.

It is especially important that knowledge graphs are easily understood by users – and not just by IT experts who build applications. As compared to basic values in a table, the connections in a knowledge graph provide a clear picture of the data. They offer highly improved accessibility, even to the deepest depths of technical data representation. Thus, experts on specific subjects can quickly and easily be involved directly in development.



Knowledge graphs are not about individual attributes but about connections and relationships, even when these are complex and diverse. This allows us to create models that can be used by humans and computers alike.



Areas of Application

INTELLIGENT CONTENT MANAGEMENT

Better, faster documentation through intelligent metadata

Personalization and the focus on the digital user experience are changing a number of fields, including technical documentation.

Today, customers do not want huge manuals that contain descriptions and every possible option, including features that may not even be relevant to the exact product they purchased. They expect operating instructions or manuals to provide support for their personal situation, in other words: for their specific product (maybe even a personalized version) and for their exact role.

This requires modularized, topic-oriented documentation and intelligent metadata that can be used to select the right topics and deliver them to the right user in the right context at the right time – whether offline, online, or integrated into the user interface of the product.

The interconnections in knowledge graphs make them the ideal technology platform for this metadata. In a knowledge graph, for example, it is perfectly sufficient to assign the component “pressure limiter” to a topic. This makes it clear that the hydraulic system is involved, and all the relevant faults and maintenance activities are then accessible without having to create links at the text level.

The intelligent information Request and Delivery Standard (iiRDS) relies on typical knowledge graph formats and thus shows what a prominent role knowledge graphs will play in the future of technical documentation.

INTELLIGENT PRODUCT INFORMATIONEN

Competitive advantages from digitalized knowledge about product relationships

Companies are increasingly moving product sales and customer support to digital channels. The major challenge in this case is creating an automated link that connects the abilities of a sales representative with the known customer situation and requirements, and with the product offering that matches those requirements.

To this end, a knowledge graph depicts products, product groups, features, use cases, target groups, and the relationships between them. On this basis, it dynamically answers the following questions:

What products are a good match for the customer's requirements?

What product features are especially important in the customer's industry?

What is the best way for the customer to use the product?

What outcome can the customer achieve with the product?

Are there restrictions that the customer has to consider with regard to how the product is used?

What cross-selling offers can we make for this product?

What accessories go with the product?

A knowledge graph is flexible enough to hold all of the product information in a central data source. With its knowledge of the relationships involved and its ability to intelligently evaluate attributes, it can determine which components are compatible with one another and then form variants dynamically – without maintaining redundant data.

It can deliver this knowledge to a variety of channels: from catalogs and web shops to product selection tools, configurators, and chatbots.



Knowledge Graphs for service

All product information in one place

Knowledge graphs bring all of the product information together in one place and reduce complexity. This creates new opportunities for customer service and after-sales.

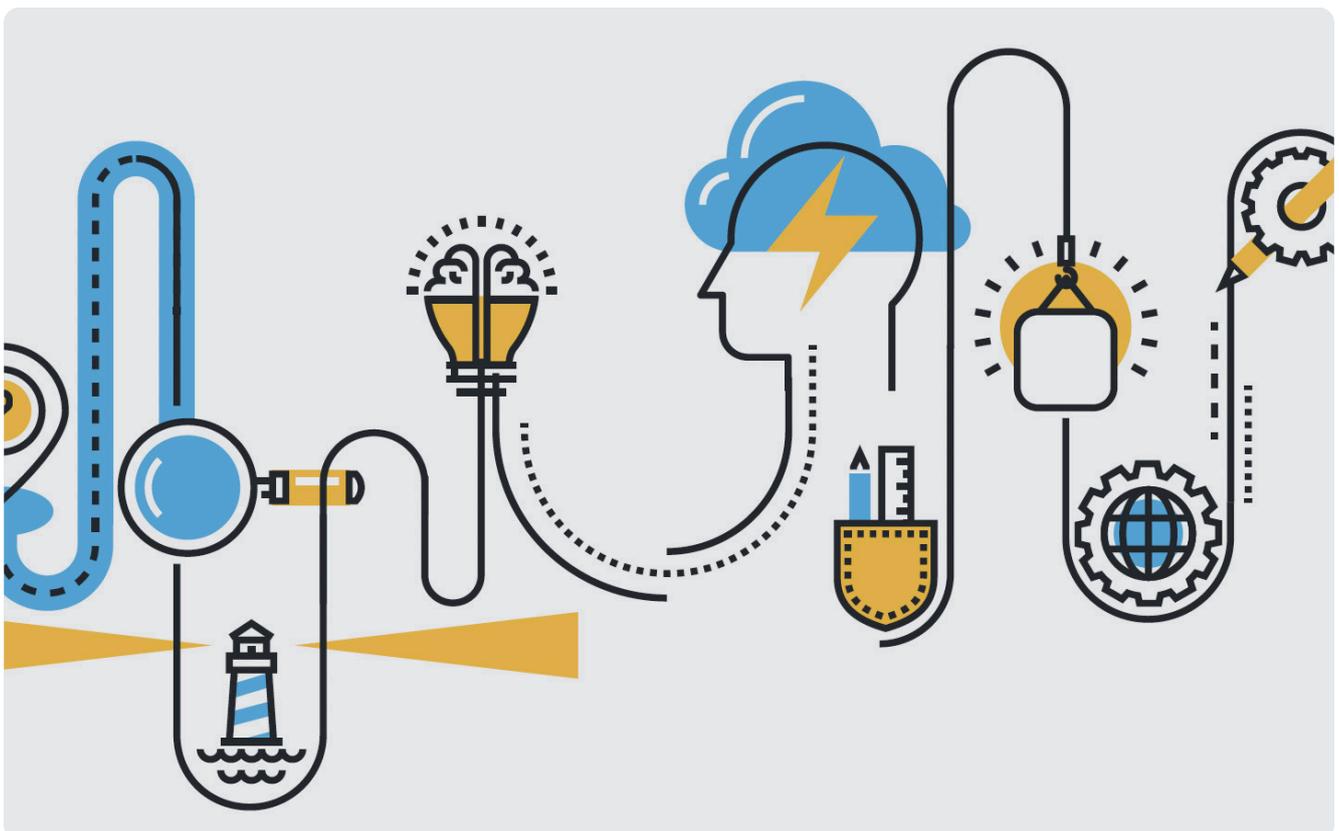
When service is needed, the knowledge graph guides the repair person or the end customer directly to the affected parts. They can then solve the problem more quickly by performing the appropriate repair or by replacing the faulty component with the right spare part. Then service staff can fall back on experience from previous cases, they can resolve even the most difficult cases without much data to go by, and they can do so in no time at all. Past experience is just what knowledge graphs provide. Based on the combination of symptoms, they can rule out certain problems more quickly or filter out solutions corresponding to the same basic faults. This saves valuable time when handling service tickets and reduces the number of escalations to the next higher service level.

Complex products, large amounts of data, expert knowledge lodged in the minds of employees – the knowledge graph brings

all of this information together. This new way of depicting relationships between products, features, and fields of use is opening up more and more applications, such as in the spare parts business or the field of self-service. This increases the potential for cross-selling and enables personalized responses to customer requirements.

Knowledge graphs overcome existing data silos with 360° access to data. All that is needed is a single piece of information to start with. With an error code, for example, a repair person or end customer can immediately identify what component is involved and determine which piece of equipment the customer uses or who the manufacturer is. This makes it easier to check inventory or ordering options.

With just a few clicks, repair staff or end customers can maneuver around silo structures without switching between applications or losing valuable time.





Capabilities of Knowledge Graphen

Now that we have familiarized ourselves with a few of the typical applications, we can take a closer look at the features needed in a knowledge graph platform to make applications successful.

Integration

Knowledge graphs are made to operate with semi-structured, unstructured, and contradictory data in complex situations – and they offer extraordinary flexibility.

These capabilities are essential, because knowledge graphs are intended to offer harmonized access to heterogeneous data sources and are rarely deployed...

Intuitive User Interfaces

When it comes to modeling the knowledge graph and configuring queries, intuitive user interfaces are important in all applications – especially if we want to use them to capture the knowledge of subject experts.

Graph-Analysis and drawing conclusions

These play a role in practically all applications. The knowledge graph follows cause-and-effect chains, for instance, or infers the suitability of a product for application based on the characteristics of a product and its features. Often, this requires the ability to identify similarities with techniques such as the weighting of graphs or rules.

Searching for and finding information

This is a primary function of knowledge graphs. Important aspects are natural language processing, i.e., the representation of linguistic knowledge (synonyms, hyponyms, multiple languages, etc.), as well as intelligent processing of search input and analysis of large quantities of text.

Visualization and interaction

Knowledge graphs represent products, service processes, and the corporate environment in a highly flexible data model. It is therefore important to have the same level of flexibility in the configuration of applications on this model, including applications with web UIs or service for external applications.

Knowledge graphs must also fulfill common requirements for production operation of a software solution in general:

Operational maturity

The use of a knowledge graph as the basis of operational applications, requires a certain "commitment," for example, in the fields of safety and deployment.

System of roles and rights

These control user access to the information in the knowledge graph.

Auditing

Logs and protocol of changes or accessing certain information.

Encrypted communication

Informationen versioning

Multilingual capability

Support of open standards

For example, JavaScript, REST, and JSON, as well as specialized metadata standards like RDFS or iiRDS.

Transaction security

For complex editing operations

Online-Backupfähigkeit

Schema-Transfer

For example, between development, test, and production systems.

Support of various platforms and deployment models

Cloud, On-Premise, support of container virtualization with Docker, etc.



How are Knowledge Graphs created?

The creation of a knowledge graph is almost always based on specific business requirements and proceeds from concrete problems and use cases. Initial ideas exist pertaining to object types and relations, thus, a data schema. This schema is flexible, however, and can be modified with little effort, at any time – which is one of the great advantages of a knowledge graph, as compared to traditional, table-based data structures.

A large part of the knowledge graph is created from the combination of existing structured data sources. The challenge is often that different systems provide redundant and contradictory information. The knowledge graph assists with the harmonization of this information with its flexibility in representing various perspectives simultaneously. Legacy systems are the leading source for the supplied data; the knowledge graph must generally reflect any changes to the source data.

As knowledge-graph-based applications are used, new objects and connections are added: e.g., new service cases, new topics in documentation, new projects, etc. Self-learning is also involved: the actions of users contain valuable implicit feedback, e.g., when users prefer certain system recommendations to others.

The graph can then be weighted differently, and new connections may be added, as necessary.

In the case of “soft data,” there is often a reversal of roles: the knowledge graph becomes the master of information. This applies, for instance, for digital workplace applications that help identify relevant projects and documents for customer needs, as well as locate products or symptoms described in service knowledge databases corresponding to a specific topic or subject. Simple, intuitive modeling is important here, along with the support of knowledge engineering.

Automation and self-learning techniques are applied here, as well. For example, we evaluate queries and keyword suggestions from users and identify candidates for new topics that can be added to the knowledge graph. These new topics are often provided from the analysis of unstructured text (in this case, we combine knowledge graphs with NLP techniques).





Summary

KNOWLEDGE GRAPHS ARE A NEW FORM OF ARTIFICIAL INTELLIGENCE, that work the way humans think and are therefore suitable for widespread use in enterprises. They are a transparent and explainable form of AI that uses its representation of knowledge to provide sound answers to complex domain-specific questions by representing subject areas and business processes in terms of their relationships. They create links and relationships, whether they be simple and uniform or complex and varied, and they enable the creation of models that can be used by both humans and computers.

KNOWLEDGE GRAPHS MAKE DISTRIBUTED KNOWLEDGE AVAILABLE AT THE RIGHT TIME AND IN THE RIGHT CONTEXT:

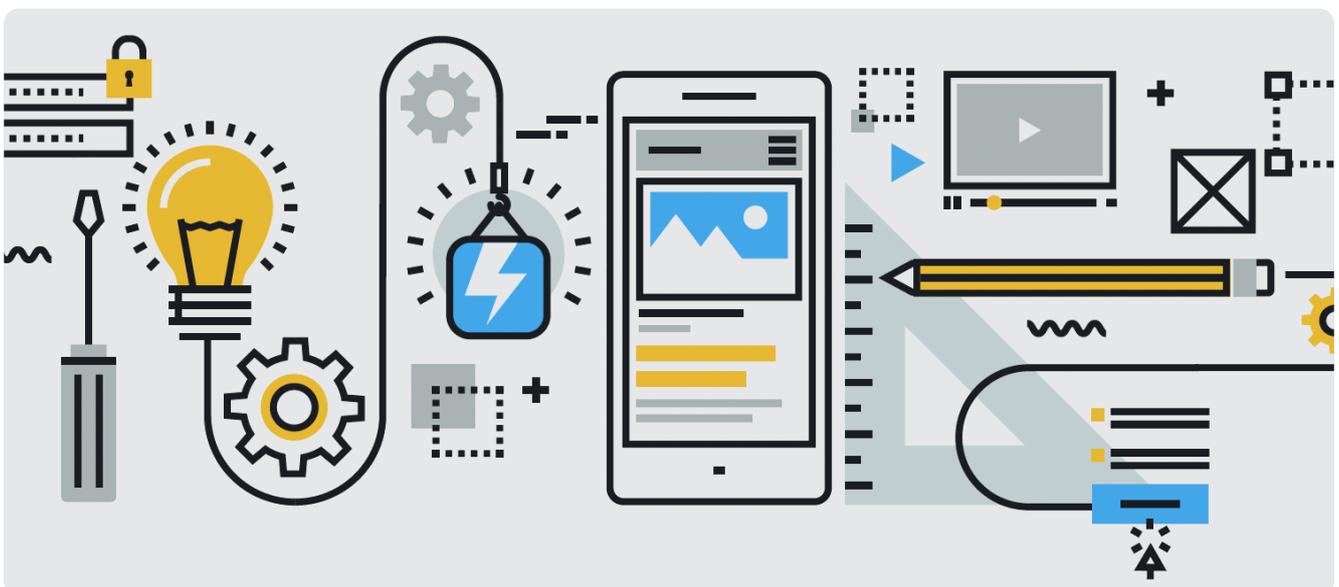
In doing so, they leverage the knowledge of the enterprise for the automation of processes, for intelligent analyses, and for assistance functions – as independent, knowledge-graph-based applications or as intelligent services for existing applications.

On this basis, fields like service, sales, after-sales, and technical documentation can be revolutionized and transformed from cost centers into profit centers.

KNOWLEDGE GRAPHS ARE IMPORTANT COMPONENTS, for adapting customer service to individual customer needs and ensuring smooth communication between internal departments. Knowledge graphs overcome data silos with 360° access to data, which allows service staff and end customers to quickly resolve even difficult cases with just a few basic parameters. This saves valuable time when handling service cases and reduces the number of escalations.

KNOWLEDGE GRAPHS ARE THE IDEAL TECHNOLOGY PLATFORM for connecting modularized, topic-oriented documentation with intelligent metadata. They ensure that users get the information they want at the right time and in the right context – offline, online, or in the user interface of their product, for example.

KNOWLEDGE GRAPHS CREATE INTELLIGENT PRODUCT INFORMATION which helps companies sell their products and provide customer support in digital channels. They reduce the complexity in service and after-sales and make it possible to tap new use cases, such as in the spare parts business or the field of self-service. This increases the potential for cross-selling and enables personalized responses to customer requirements.





Contact us!

Empolis solutions enable companies and organizations to analyze, interpret, and automatically process exponentially growing quantities of structured and unstructured data. This allows them to use their knowledge capital to optimize essential business processes. Accordingly, decision makers, employees, and customers receive exactly the information appropriate to the situation or task, which allows them to make better decisions more quickly.



Find more information on www.empolis.com